

Changepoint Detection over Graphs with the Spectral Scan Statistic

James Sharpnack ^{*1,2}, Alessandro Rinaldo ^{†2}, and Aarti Singh ^{‡ 1}

¹Machine Learning Department, Carnegie Mellon University

²Statistics Department, Carnegie Mellon University

March 6, 2013

Abstract

We consider the change-point detection problem of deciding, based on noisy measurements, whether an unknown signal over a given graph is constant or is instead piecewise constant over two connected induced subgraphs of relatively low cut size. We analyze the corresponding generalized likelihood ratio (GLR) statistics and relate it to the problem of finding a sparsest cut in a graph. We develop a tractable relaxation of the GLR statistic based on the combinatorial Laplacian of the graph, which we call the spectral scan statistic, and analyze its properties. We show how its performance as a testing procedure depends directly on the spectrum of the graph, and use this result to explicitly derive its asymptotic properties on few significant graph topologies. Finally, we demonstrate both theoretically and by simulations that the spectral scan statistic can outperform naive testing procedures based on edge thresholding and χ^2 testing.

1 Introduction

In this article we are concerned with the basic but fundamental task of deciding whether a given graph, over which a noisy signal is observed, contains a cluster of anomalous or activated nodes comprising an induced connected subgraph. Such a problem is highly relevant in a variety of scientific areas, such as community detection in social networks, surveillance, disease outbreak detection, biomedical imaging, sensor network detection, gene network analysis, environmental monitoring and malware detection. Recent theoretical contributions in the statistical literature (see, e.g., Arias-Castro et al. [2005, 2008, 2011], Addario-Berry et al. [2010]) have detailed the inherent difficulty of such a testing problem in relatively simplified settings and under specific conditions on the graph topology. From a practical standpoint, the natural algorithm for detection of anomalous clusters of activity in graphs is the the generalized likelihood ratio test (GLRT) or scan statistic, a computationally intensive procedure that entails scanning all well connected clusters and testing individually for anomalous activation. Unfortunately, its performance over general graphs is not well understood, and little attention has been paid to determining alternative, computationally tractable, procedures.

In this article we assume that the class of clusters of activation consists of sub-graphs of small cut size. We believe this is a natural and realistic assumption which, as we demonstrate below, allows us to explicitly incorporate into the detection problem the properties of the graph topology through its spectrum. In particular, we show that the GLRT is an integer program with a term in the objective that corresponds to the sparsest cut in a graph, a known NP-hard problem. With this in mind, we propose a relaxation of the GLRT, called the spectral scan statistic, which is based on the combinatorial Laplacian of the graph and, importantly, is a tractable program. As our main result, we derive theoretical guarantees for the performance of the spectral scan statistic, which hold for any graph and are based on the spectral measure of the combinatorial Laplacian. For comparison purposes, we derive theoretical guarantees for

*jsharpna@cs.cmu.edu

†arinaldo@cmu.edu

‡aarti@cs.cmu.edu

two simple estimators, the edge thresholding and the χ^2 test. We conclude our study by applying the main result to balanced binary trees, the lattice, and Kronecker graphs, giving us precise asymptotic results. We find that, modulo logarithm terms, the spectral scan statistic has nearly optimal power for balanced binary trees. Simulations for these models verify that the spectral scan statistic dominates the simple estimators.

Contributions. Our contributions are as follows. (1) We define a new class of activation patterns based on the notion of small cut size that reflects in a natural way the topological properties of the graph. (2) We analyze the corresponding GLR statistics and show that it is indeed related to the problem of finding sparsest cuts. We then develop a computationally tractable relaxation of the GLR statistic, called the spectral scan statistic and analyze its properties. In our main theoretical result, we show that the performance of the spectral scan statistic depends explicitly on the spectral properties of the graph. (3) Using such results we are able to characterize in a very explicit form the performance of the spectral scan statistic on a few notable graph topologies and demonstrate its superiority over naive detectors, such as the edge thresholding and the χ^2 test. (4) Finally, we have formulated the detection problem under more general and realistic scenarios, which involve composite null and alternative hypotheses as opposed to simple hypotheses as is customary in the theoretical statistical literature on this subject.

Related Work. Normal means testing in high-dimensions is a well established and fundamental problem in statistics (see, e.g., Ingster and Suslina [2003]). A significant portion of the recent work in this area (Arias-Castro et al. [2005, 2008, 2011], Addario-Berry et al. [2010]) has focused on incorporating structural assumptions on the signal, as a way to mitigate the effect of high-dimensionality and also because many real-life problems can be represented as instances of the normal means problem with graph-structured signals (see, for an example, Jacob et al. [2010]). These contributions have considered the generalized likelihood ratio test of means when the alternative hypothesis takes on the form of a combinatorial space. However, the performance of such test has been analyzed only for certain types of graphs, and it is unclear to what extent those analyses extend to general graph topologies. Moreover, while much is known about the theoretical performance of the GLRT, no mention is made about its computational feasibility. Another line of research relevant to our problem is the optimal fail detection with nuisance parameters and matched subspace detection in the signal processing literature: see, e.g. Scharf and Friedlander [1994], Baygün and Hero [1995], Fouladirad and Nikiforov [2005], Fouladirad et al. [2008]. Though our problem can be cast as a special case of the more general problem of optimal testing of a linear subspace under nuisance parameters considered in that line of work, the focus on a graph-structured signal, as well as the type of analysis based on the interplay between the scan statistics and the spectral properties of the graph contained in our work, are novel.

1.1 Problem Setup

We now formalize the problem of detecting a change of signal over the vertices of a graph from noisy observations in the high-dimensional setting. For a given connected, undirected, possibly weighted graph $G = (V, E)$ on $|V| = n$ nodes, we observe *one* realization of the random vector

$$\mathbf{y} = \boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (1)$$

where $\boldsymbol{\beta} \in \mathbb{R}^V$ and $\boldsymbol{\epsilon} \sim N(0, \sigma^2 \mathbf{I}_n)$, with σ^2 known. We will assume that there are two groups of constant activation for the signal $\boldsymbol{\beta}$, namely that there exists a subset $C \subset V$ such that $\boldsymbol{\beta}$ is constant within both C and its complement $\bar{C} = V \setminus C$. We formalize this assumption by writing

$$\boldsymbol{\beta} = \mu \mathbf{1} + \delta \mathbf{1}_C, \quad (2)$$

where $\mu, \delta \in \mathbb{R}$ are unknown parameters, $\mathbf{1} \in \mathbb{R}^V$ is a n -dimensional vector of ones and $\mathbf{1}_C$ is the indicator function of the subset C . The parameter μ can be thought of as the magnitude of the background signal and is a nuisance parameter, while δ quantifies the gap in signal between the two clusters. Setting $\bar{\boldsymbol{\beta}} = \mathbf{1}^\top \boldsymbol{\beta} / n$, we will use $\|\boldsymbol{\beta} - \bar{\boldsymbol{\beta}}\|$ to measure the energy of the signal (note that this quantity is independent of μ), and we will define the signal-to-noise ratio (SNR) to be

$$\frac{\|\boldsymbol{\beta} - \bar{\boldsymbol{\beta}}\|}{\sigma} = \sqrt{\frac{|C||\bar{C}|}{n}} \frac{\delta}{\sigma}.$$

We will not assume any knowledge of the true clustering (C, \bar{C}) , other than that it belongs to a given class \mathcal{C} of bipartitions (C, \bar{C}) of V such that C and \bar{C} are both large and can be easily disconnected, in that they have low cut size. Formally, we define, for some $\rho > 0$,

$$\mathcal{C} = \mathcal{C}(\rho) = \left\{ C \subset V, C \neq \emptyset : \frac{|\partial C|}{|C||\bar{C}|} \leq \frac{\rho}{|V|} \right\}, \quad (3)$$

where $\partial C = \{(i, j) \in E : i \in C, j \in \bar{C}\}$ is the boundary of C . Note that \mathcal{C} is a symmetric class in the sense that $C \in \mathcal{C}$ if and only if $\bar{C} \in \mathcal{C}$. We are interested in the problem of testing whether the gap parameter δ in equation (2) is zero (i.e. the signal β is constant) or it is non-zero for some $C \in \mathcal{C}$, regardless of the value of μ . Thus, we can naturally cast our structured change-point detection problem as the following composite hypothesis testing problem:

$$H_0: \beta \in \Theta_0 \quad \text{vs} \quad H_1: \beta \in \Theta_1, \quad (4)$$

where $\Theta_0 = \{\mu \mathbf{1}, \mu \in \mathbb{R}\}$ and $\Theta_1 = \{\mathbf{1}_C \mu + \mathbf{1}_C \delta, \mu \in \mathbb{R}, \delta \in \mathbb{R} \setminus \{0\}, C \in \mathcal{C}\}$. Notice that the alternative can be written as the join over \mathcal{C} of disjoint composite alternatives of the form $H_1^C: \beta \in \Theta_1^C := \{\mathbf{1}_C \mu + \mathbf{1}_C \delta, \mu \in \mathbb{R}, \delta \in \mathbb{R} \setminus \{0\}\}$, $C \in \mathcal{C}$.

To make our analysis meaningful, we measure the difficulty of the detection problem in terms of the energy parameter by assuming that, for some $\eta > 0$, $\|\beta - \bar{\beta}\| > \eta$, $\forall \beta \in \Theta_1$. Thus, we can think of η as the minimal degree of separation between the null and alternative hypotheses. Below we will analyze asymptotic conditions under which the hypothesis testing problem described above is feasible, in a sense made precise in the next definition, when the size of the graph n increases unboundedly. To this end, we will further assume that the relevant parameters of the model, η , σ , δ and ρ change with n as well, even though we will not make such dependence explicit in our notation for ease of readability. Our results establish conditions for asymptotic disinguishability as a function of the SNR η/σ and ρ and the spectrum of the graph G .

Definition 1. Let P_θ denote the distribution of \mathbf{y} induced by the model (1), where $\theta \in \Theta_0 \cup \Theta_1$. For a given statistic $S(\mathbf{y})$ and threshold $\tau \in \mathbb{R}$, let $T = T(\mathbf{y})$ be 1 if $S(\mathbf{y}) > \tau$ and 0 otherwise. We say that the hypotheses H_0 and H_1 are **asymptotically distinguished by the test T** if

$$\sup_{\theta \in H_0} \mathbb{P}_\theta\{T = 1\} \rightarrow 0 \quad \text{and} \quad \sup_{\theta \in H_1} \mathbb{P}_\theta\{T = 0\} \rightarrow 0, \quad (5)$$

where the limit is taken as $n \rightarrow \infty$. We say that H_0 and H_1 are **asymptotically indistinguishable** if there does not exist any test for which the above limits hold.

Notation. We will need some mathematical terminology from algebraic graph theory (Godsil et al. [2001]). A central object to our analysis is the *combinatorial Laplacian* matrix $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where $\mathbf{W} = (I\{(v, w) \in E\})_{v, w \in V}$ is the adjacency matrix of the graph G and $\mathbf{D} = \text{diag}\{d_v\}_{v \in V}$ is the diagonal matrix of node degrees, $d_v = \sum_{w \in V} W_{v, w}$, $v \in V$. If the graph is weighted then $W_{v, w}$ reflects this. We will denote the eigenvalues of \mathbf{L} with $\{\lambda_i\}_{i=1}^n$, which we will always take in increasing order. Since G is connected, the smaller eigenvalue $\lambda_1 = 0$, with corresponding eigenvector, $\mathbf{1}$. λ_2 is known as the *algebraic connectivity* and is lower bounded by $4[n \text{diam}(G)]^{-1}$ where $\text{diam}(G)$ is the diameter of the graph. Throughout this study we use Bachmann-Landau notation for asymptotic statements: if $a_n/b_n \rightarrow 0$ then $a_n = o(b_n)$ and $b_n = \omega(a_n)$.

2 Methods

The hypothesis testing problem at hand presents two challenges: (1) the model contains an unbounded nuisance parameter $\mu \in \mathbb{R}$ and (2) the alternative hypothesis is comprised of a finite disjoint union of composite hypotheses indexed by \mathcal{C} . These features set our problem apart from virtually all existing work of structured normal means problems (see, e.g. Arias-Castro et al. [2005, 2008, 2011], Addario-Berry et al. [2010]), which does not consider nuisance parameters and relies on a simplified framework consisting of a simple null hypothesis and a composite hypothesis consisting of disjoint unions of simple alternatives. Having nuisance parameters and composite hypothesis require a more sophisticated analysis.

We will eliminate the interference caused by the nuisance parameter by considering test procedures that are independent of μ . The formal justification for this choice is based on the theory of optimal invariant hypothesis testing (see, e.g., Lehmann and Romano [2005]) and of uniformly best constant power tests (see Wald [1943]). Due to space limitations we will not provide the details and refer the reader to Fouladirad et al. [2008], Fouladirad and Nikiforov [2005], Fillatre and Nikiforov [2007], Fillatre [2012], Scharf and Friedlander [1994], Baygün and Hero [1995] and references therein for in depth-treatments of these issues related to the model a hand.

For the simpler problem of testing H_0 versus H_1^C for some $C \subset V$, the optimal test is based on the likelihood ratio (LR) statistic (see the proof of Lemma 2 below for a derivation)

$$2 \log \Lambda_C(\mathbf{y}) = \log \left(\frac{\sup_{\theta \in \Theta_1} f_\theta(\mathbf{y})}{\sup_{\theta \in \Theta_0} f_\theta(\mathbf{y})} \right) = \frac{1}{\sigma^2} \frac{|V|}{|C||\bar{C}|} \left(\sum_{v \in C} \tilde{\mathbf{y}}_v \right)^2, \quad (6)$$

where $\tilde{\mathbf{y}} = \mathbf{y} - \bar{\mathbf{y}} = (\tilde{\mathbf{y}}_v, v \in V)$ and f_θ is the Lebesgue density of P_θ . This test rejects H_0 for large values of $\Lambda_C(\mathbf{y})$. Optimality follows from the fact that the statistical model we consider has the monotone likelihood ratio property.

When testing against composite alternatives, like in our case, it is customary to consider instead the generalized likelihood ratio (GLR) statistic, which in our case reduces to

$$\hat{g} = \max_{C \in \mathcal{C}(\rho)} 2\sigma^2 \log \Lambda_C(\mathbf{y}).$$

Through manipulations of the likelihoods, we find that the GLR statistic has a very convenient form which is tied to the spectral properties of the graph G via its Laplacian.

Lemma 2. *Let $\tilde{\mathbf{y}} = \mathbf{y} - \mathbf{1}(\frac{1}{n} \sum_{v \in V} \mathbf{y}_v)$ and $\mathbf{K} = \mathbf{I} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top$. Then*

$$\hat{g} = \max_{\mathbf{x} \in \{0,1\}^n} \frac{\mathbf{x}^\top \tilde{\mathbf{y}} \tilde{\mathbf{y}}^\top \mathbf{x}}{\mathbf{x}^\top \mathbf{K} \mathbf{x}} \text{ s.t. } \frac{\mathbf{x}^\top \mathbf{L} \mathbf{x}}{\mathbf{x}^\top \mathbf{K} \mathbf{x}} \leq \rho, \quad (7)$$

where \mathbf{L} is the combinatorial Laplacian of the graph G .

The proof is provided in the appendix. The savvy reader will notice the connection between (7) and the graph sparsest cut program. By Lagrangian duality, we see that the program (7) is equivalent to (for some Lagrangian parameter ν)

$$\min_{C \subseteq V} \frac{|\partial C|}{|C||\bar{C}|} - \nu \frac{(\sum_{i \in C} \tilde{y}_i)^2}{|C||\bar{C}|}$$

the first term of which is precisely the *sparsest cut* objective, and the second term drives the solution C to have positive within cluster empirical correlations. The sparsest cut program is known to be NP-hard, with poly-time algorithms known for trees and planar graphs (Matula and Shahrokhi [1990]). Because of this fact, approximate algorithms have been proposed over the past two decades, most notably the uniform multicommodity flow approach of (Leighton and Rao [1988], Shmoys [1997]) and the semi-definite relaxation of the cut metric (Arora et al. [2009]). Hagen and Kahng [1992] observed that the minimum cut sparsity is bounded by the algebraic connectivity (λ_2), suggesting the Fiedler vector (i.e. the second eigenvector of \mathbf{L}) to be an appropriate relaxation of the characteristic vector of the cut. Moreover, the well known Cheeger inequality shows that the minimum cut sparsity (in a regular graph) is bounded by the algebraic connectivity (see Chung [2004]). We will follow the tradition of bounding sparsity with the algebraic connectivity, and provide a surrogate estimator to the scan statistic based on this simple spectral relaxation.

Proposition 3. *Define the Spectral Scan Statistic (SSS) as*

$$\hat{s} = \sup_{\mathbf{x} \in \mathbb{R}^n} (\mathbf{x}^\top \tilde{\mathbf{y}})^2 \text{ s.t. } \mathbf{x}^\top \mathbf{L} \mathbf{x} \leq \rho, \|\mathbf{x}\| \leq 1, \mathbf{x}^\top \mathbf{1} = 0.$$

Then the GLR statistic is bounded by the SSS: $\hat{g} \leq \hat{s}$.

Proof. First let us notice that $\mathbf{K} = \mathbf{I} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top$ is the projection onto the subspace orthogonal to $\mathbf{1}$. Because \mathbf{K} is thus idempotent, $\tilde{\mathbf{y}}\mathbf{1} = 0$, and $\mathbf{L}\mathbf{1} = 0$ we can rewrite

$$\hat{g} = \max_{\mathbf{x} \in \{0,1\}^n \setminus \{\mathbf{0}, \mathbf{1}\}} \frac{(\mathbf{K}\mathbf{x})^\top \tilde{\mathbf{y}} \tilde{\mathbf{y}}^\top (\mathbf{K}\mathbf{x})}{(\mathbf{K}\mathbf{x})^\top (\mathbf{K}\mathbf{x})} \text{ s.t. } \frac{(\mathbf{K}\mathbf{x})^\top \mathbf{L} (\mathbf{K}\mathbf{x})}{(\mathbf{K}\mathbf{x})^\top (\mathbf{K}\mathbf{x})} \leq \rho$$

So, we have the following relaxation,

$$\hat{g} \leq \max_{\mathbf{x} \neq \mathbf{0}, \mathbf{x}^\top \mathbf{1} = 0} \frac{\mathbf{x}^\top \tilde{\mathbf{y}} \tilde{\mathbf{y}}^\top \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} \text{ s.t. } \frac{\mathbf{x}^\top \mathbf{L} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} \leq \rho = \hat{s}$$

□

Remark 4. By Lagrangian duality and the Courant-Fischer theorem, the spectral scan statistic can be written as

$$\hat{s} = \min_{\nu > 0} \chi(\tilde{\mathbf{y}}\tilde{\mathbf{y}}^\top - \nu\Delta) + \nu\rho$$

where $\chi(A)$ is the maximum non-zero eigenvalue of the matrix A .

Notice that because the domain $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\top \mathbf{L}\mathbf{x} \leq \rho, \|\mathbf{x}\| \leq 1, \mathbf{x}^\top \mathbf{1} = 0\}$ is symmetric around the origin, this is precisely the square of the solution to

$$\sqrt{\hat{s}} = \sup_{\mathbf{x} \in \mathbb{R}^n} \mathbf{x}^\top \mathbf{y} \text{ s.t. } \mathbf{x}^\top \mathbf{L}\mathbf{x} \leq \rho, \|\mathbf{x}\| \leq 1, \mathbf{x}^\top \mathbf{1} = 0, \quad (8)$$

where we have used the fact that $\mathbf{x}^\top \tilde{\mathbf{y}} = ((\mathbf{I} - \frac{1}{n}\mathbf{1}\mathbf{1}^\top)\mathbf{x})^\top \mathbf{y} = \mathbf{x}^\top \mathbf{y}$ because $\mathbf{x}^\top \mathbf{1} = 0$ within \mathcal{X} . This previous formulation shows that the SSS is related to the supremum of a Gaussian process over \mathcal{X} . This fact will turn out to be extremely convenient, as we show next.

3 Theoretical Analysis

We first derive a simple condition for asymptotic indistinguishability based on testing the null versus a single component in the alternative. A more refined analysis of the lower bound for the general hypothesis (4) is beyond the scope of this article.

Theorem 5. Suppose that there exists $C \in \mathcal{C}$ such that $\frac{|\bar{C}|}{|C|} \asymp 1$. Then H_0 and H_1 are asymptotically indistinguishable if $\eta/\sigma = o(1)$.

The proof is in the appendix. We will analyze the performance of the SSS statistic by relying on its representation (8) as the square of the supremum of a Gaussian process. We draw heavily on the theory of the generic chaining, perfected in Talagrand [2005], which essentially reduces the problem of computing bounds on the expected supremum of Gaussian processes to geometric properties of its index space. Recall that, under alternative hypothesis, $\|\beta - \bar{\beta}\| \geq \eta$ uniformly over Θ_1 .

Theorem 6. The following hold with probability at least $1 - \delta$. Under the null H_0

$$\hat{s} \leq \left(\sqrt{2\sigma^2 \sum_{i>1} \min\{1, \rho\lambda_i^{-1}\}} + \sqrt{2\sigma^2 \log \frac{2}{\delta}} \right)^2,$$

while the alternative H_1

$$\hat{s} \geq \left(\eta - \sqrt{2\sigma^2 \log \frac{2}{\delta}} \right)^2.$$

Proof. We use generic chaining to control the process $\{\mathbf{x}^\top \mathbf{y}\}_{\mathbf{x} \in \mathcal{X}}$ appearing in the SSS. First, we notice that the index set \mathcal{X} is the intersection of an ellipsoid and the unit ball, which is the intuition behind the following lemma.

Lemma 7. Let \mathbf{L} have spectrum $\{\lambda_i\}_{i=1}^n$. Then under H_0 ,

$$\mathbb{E} \sup_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \mathbf{y} \leq \sqrt{2\sigma^2 \sum_{i>1} \min\{1, \rho\lambda_i^{-1}\}}.$$

The proof is provided in the appendix. We then can use the well known phenomena, that the supremum of a Gaussian process concentrates around its expectation (see the appendix). Hence, by Lemma 14 the first statement in Theorem 6 holds. The second statement follows by applying standard concentration results to the univariate Gaussian $\frac{\beta - \bar{\beta}}{\|\beta - \bar{\beta}\|}^\top \mathbf{y}$ and noticing that $\frac{\beta - \bar{\beta}}{\|\beta - \bar{\beta}\|} \in \mathcal{X}$ and $\mathbb{E} \frac{(\beta - \bar{\beta})^\top}{\|\beta - \bar{\beta}\|} \mathbf{y} = \|\beta - \bar{\beta}\| \geq \eta$ under H_1 . \square

As a corollary we will provide sufficient conditions for asymptotic distinguishability that depend on the spectrum of the Laplacian \mathbf{L} . As we will show in the next section, these conditions can be applied to a number of graph topologies whose spectral properties are known.

Corollary 8. *The null and alternative, as described in Thm. 6, are asymptotically distinguished by \hat{s} and $g_\nu(\mathbf{y})$ if*

$$\frac{\eta}{\sigma} = \omega \left(\sqrt{\sum_{i>1} \min\{1, \rho\lambda_i^{-1}\}} \right) \quad (9)$$

Other stronger sufficient conditions are

$$\frac{\eta}{\sigma} = \omega \left(\sqrt{k + \frac{(n-k)\rho}{\lambda_{k+1}}} \right) \quad (10)$$

if k is large enough that $\lambda_{k+1} > \rho$.

Proof. To see equation (9) we note that, due to Theorem 6, if

$$\sqrt{2\sigma^2 \sum_{i>1} \min\{1, \rho\lambda_i^{-1}\}} + \sqrt{2\sigma^2 \log \frac{2}{\delta}} = o \left(\eta - \sqrt{2\sigma^2 \log \frac{2}{\delta}} \right)$$

then we attain asymptotic distinguishability by choosing any threshold τ between, and sufficiently far from, the left and right hand side of the previous display. To show equation (10) we note that by choosing k such that $\lambda_{k+1} > \rho$ we see that

$$\sum_{1 < i \leq k} \min\{1, \rho\lambda_i^{-1}\} \leq k \Rightarrow \sum_{i > k} \min\{1, \rho\lambda_i^{-1}\} \leq (n-k) \frac{\rho}{\lambda_{k+1}}.$$

□

Interestingly, there are no logarithmic terms in (9) that usually accompany uniform bounds of this type, which is attributed to the generic chaining. Notice that the left hand side of (9) is always less than $\sqrt{n-1}$, which we will see characterizes the performance of the naive estimator $\|\tilde{\mathbf{y}}\|$.

For comparison, we consider the performance of two naive procedure for detection: the energy detector, which reject H_0 if $\|\tilde{\mathbf{y}}\|^2$ is too large and the edge thresholding detector, which reject H_0 if $\max_{(v,w) \in E} |\mathbf{y}_v - \mathbf{y}_w|$ is large.

Theorem 9. *H_0 and H_1 are asymptotically distinguished by $\|\tilde{\mathbf{y}}\|$ if and only if*

$$\frac{\eta}{\sigma} = \omega(\sqrt{n-1}).$$

The proof (given in the appendix) is a standard χ^2 analysis. In Sharpnack et al. [2012] the authors examined the problem of exact recovery of cluster boundaries in the graph-structured normal means problem by taking differences between observations corresponding to adjacent nodes. The following result stems from Theorem 2.1 of Sharpnack et al. [2012], and the fact that $|C||\bar{C}|/n$ scales like $\min\{|C|, |\bar{C}|\}$ up to a factor of 2.

Theorem 10. *H_0 and H_1 are asymptotically distinguished by $\max_{(v,w) \in E} |\mathbf{y}_v - \mathbf{y}_w|$ if*

$$\frac{\eta}{\sigma} = \omega \left(\sqrt{\max_{C \in \mathcal{C}, |C| \leq n/2} |C| \log n} \right).$$

If \mathcal{C} contains balanced clusters, i.e. bipartitions (C, \bar{C}) such that $\frac{|C|}{|\bar{C}|} \asymp 1$, then this result matches the scaling in Theorem 9 up to a log factor.

4 Specific Graph Models

In this section we demonstrate the power and flexibility of Theorem 6 by analyzing in detail the performance of the spectral scan statistic over three important graph topologies: balanced binary trees, the s -dimensional lattice and the Kronecker graphs (see Leskovec and Faloutsos [2007], Leskovec et al. [2010]).

4.1 Balanced Binary Trees

We begin the analysis of the spectral scan statistic by applying it to the balanced binary tree (BBT) of depth ℓ . The class of signals that we will consider have clusters of constant signal which are subtrees of size at least cn^α for $0 < c \leq 1/2, 0 < \alpha \leq 1$. Hence, the cut size of the signals are 1 and $\rho = [cn^\alpha(1 - cn^{\alpha-1})]^{-1}$.

Corollary 11. *For the balanced binary tree with n vertices, the spectral scan statistic can asymptotically distinguish H_0 from signals with $\rho = n[cn^\alpha(n - cn^\alpha)]^{-1}$ if the SNR is stronger than*

$$\frac{\eta}{\sigma} = \omega(n^{\frac{1-\alpha}{2}} \log n).$$

We simulate the probability of correct discovery of change-points (rejecting H_0 when the truth is H_1) versus the probability of false alarm (falsely rejecting H_0). These are given for the four estimators in Figure 1 and for the SSS as $n = 2^{\ell+1} - 1$ increases. In these simulations a subtree at level 2 (of size $n/4$) was chosen as C , the gap-to-noise ratio is fixed at $\delta/\sigma = 0.8$, and $\rho = 4/n$. We see that even in the low n regime, exploiting the graph structure is essential to improve the power of testing H_0 against H_1 . As n increases with δ/σ fixed the performance of the SSS dramatically increases.

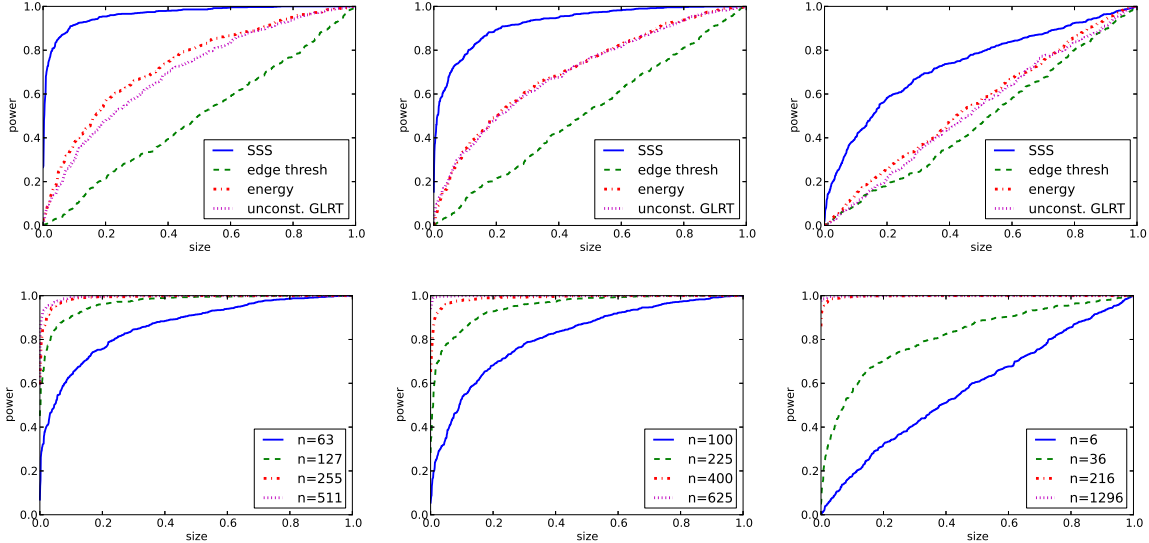


Figure 1: Above: the simulated probability of correct discovery (power) against false alarm (size) of the SSS compared to the energy detector, edge thresholding and the unconstrained GLRT of the BBT (left), Lattice (middle), and Kronecker graph (right). Below: the performance as n increases.

4.2 Lattice

We will analyze the performance guarantees of the SSS over the 2-dimensional lattice graph with p vertices along each dimension ($n = p^2$). We will assume that $\rho = Cn^{-1/2}$, as this is the cut sparsity of rectangles that have a low surface area to volume ratio. By a simple Fourier analysis (see Sharpnack and Singh [2010]), we know that the Laplacian eigenvalues are $2(2 - \cos(2\pi i_1/p) - \cos(2\pi i_2/p))$ for all $i_1, i_2 \in [p]$. We will appeal to (10). Because $1 - \cos(2\pi i_1/p) \approx (2\pi i_1/p)^2$ for $i_1 \ll p$, if we rewrite $i = (i_1, i_2)$ for $i_1, i_2 \in [p]$ then $\lambda_{(i_1, i_2)} \approx \frac{8\pi^2}{n}(i_1^2 + i_2^2)$. Hence,

$$k \approx |\{(i_1, i_2) : i_1^2 + i_2^2 \leq \frac{n}{8\pi^2} \lambda_{k+1}\}| \leq |\{i_1 : i_1^2 \leq \frac{n}{8\pi^2} \lambda_{k+1}\}|^2 = \lceil \frac{n}{8\pi^2} \lambda_{k+1} \rceil$$

Then by choosing $\lambda_{k+1} \asymp \sqrt{\rho}$ the term in the root of the LHS of (10) is bounded by, $\lceil \frac{n}{8\pi^2} \lambda_{k+1} \rceil + \frac{\rho n}{\lambda_{k+1}} \asymp n\sqrt{\rho} \asymp n^{3/4}$ modulo lower order terms. We arrive at the following conclusion,

Corollary 12. *For the $p \times p$ square lattice, the spectral scan statistic can asymptotically distinguish H_0 from signals with cut size $Cn^{-1/2}$ if the SNR is stronger than,*

$$\frac{\eta}{\sigma} = \omega(n^{3/8})$$

We demonstrate the improvement of the SSS over competing tests in Figure 1. In these simulations a $\sqrt{n}/2 \times \sqrt{n}/2$ square was chosen to be C with $\rho = 4/\sqrt{n}$. Despite the weaker guarantee in Corollary 12 the SSS demonstrates the importance of exploiting the graph structure.

4.3 Kronecker Graphs

Much of the research in complex networks has focused on observing statistical phenomena that is common across many data sources. The most notable of these are that the degree distribution obeys a power law (Faloutsos et al. [1999]) and networks are often found to have small diameter (Milgram [1967]). A class of graphs that satisfy these, while providing a simple modelling platform are the Kronecker graphs (see Leskovec and Faloutsos [2007], Leskovec et al. [2010]). Let H_1 and H_2 be graphs on p vertices with Laplacians $\mathbf{L}_1, \mathbf{L}_2$ and edge sets E_1, E_2 respectively. The Kronecker product, $H_1 \otimes H_2$, is the graph over vertices $[p] \times [p]$ such that there is an edge $((i_1, i_2), (j_1, j_2))$ if $i_1 = j_1$ and $(i_2, j_2) \in E_2$ or $i_2 = j_2$ and $(i_1, j_1) \in E_1$. We will construct graphs that have a multi-scale topology using the Kronecker product. Let the multiplication of a graph by a scalar indicate that we multiply each edge weight by that scalar. First let H be a connected graph with p vertices. Then the graph G for $\ell > 0$ levels is defined as

$$\frac{1}{p^{\ell-1}}H \otimes \frac{1}{p^{\ell-2}}H \otimes \dots \otimes \frac{1}{p}H \otimes H$$

The choice of multipliers ensures that it is easier to make cuts at the more coarse scale. Notice that all of the previous results have held for weighted graphs.

Corollary 13. *For G be the Kronecker product graph described above with $n = p^\ell$ vertices, the spectral scan statistic can asymptotically distinguish H_0 from signals with cuts within the k coarsest scale ($\rho \propto p^{2k-\ell-1}$), if the SNR is stronger than,*

$$\frac{\eta}{\sigma} = \omega(p^2(\ell + 2)n^{(2k+1)/\ell})$$

The proof and an explanation of ρ is in the appendix. Again, we demonstrate the improvement of the SSS over competing tests in Figure 1. For these simulations the base graph H was chosen to be two triangles (K_3) connected by a single edge ($p = 6$). At the coarsest scale one of the K_3 subgraphs was chosen to be C with $\rho = 4/n$.

5 Discussion

We studied the heretofore unaddressed problem of how to tractably detect change-points in networks under Gaussian noise. To this end we developed the spectral scan statistic, suggesting it as a computationally feasible alternative to the GLRT. We completely characterized the performance of the SSS for any graph in terms of the spectrum of the combinatorial Laplacian. For comparison purposes, we developed theoretical guarantees for two simple estimators. We applied the main result to three graph models: binary balanced trees, the lattice and Kronecker graph. We see that not only is it statistically inadmissible to ignore graph structure, but for the balanced tree the SSS gives near optimal performance. This claim is backed by both simulation and theory.

Acknowledgements

This research is supported in part by AFOSR under grant FA9550-10-1-0382 and NSF under grant IIS-1116458.

References

L. Addario-Berry, N. Broutin, L. Devroye, and G. Lugosi. On combinatorial testing problems. *The Annals of Statistics*, 38(5):3063–3092, 2010.

- E. Arias-Castro, D. Donoho, and X. Huo. Near-optimal detection of geometric objects by fast multiscale methods. *IEEE Trans. Inform. Theory*, 51(7):2402–2425, 2005.
- E. Arias-Castro, E.J. Candes, H. Helgason, and O. Zeitouni. Searching for a trail of evidence in a maze. *The Annals of Statistics*, 36(4):1726–1757, 2008.
- E. Arias-Castro, E.J. Candes, and A. Durand. Detection of an anomalous cluster in a network. *The Annals of Statistics*, 39(1):278–304, 2011.
- S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. *Journal of the ACM (JACM)*, 56(2):5, 2009.
- M. Azizyan and A. Singh. Subspace detection of high-dimensional vectors using compressive sampling. *IEEE SSP*, 2012.
- B. Baygün and A. O. Hero. Optimal simultaneous detection and estimation under a false alarm constraint. *Signal Processing, IEEE Transactions on*, 41(3):688–703, 1995.
- C. Borell. The brunn-minkowski inequality in gauss space. *Inventiones Mathematicae*, 30(2):207–216, 1975.
- F. Chung. Discrete isoperimetric inequalities. *Surveys in Differential Geometry IX, International Press*, pages 53–82, 2004.
- B. Cirelson, I. Ibragimov, and V. Sudakov. Norms of gaussian sample functions. In *Proceedings of the Third JapanUSSR Symposium on Probability Theory*, pages 20–41. Springer, 1976.
- M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM Computer Communication Review*, volume 29, pages 251–262. ACM, 1999.
- M. Fiedler. Eigenvectors of acyclic matrices. *Czechoslovak Mathematical Journal*, 25(4):607–618, 1975.
- L. Fillatre. Asymptotically uniformly minimax detection and isolation in network monitoring. to appear in *Signal Processing, IEEE Transactions on*, 2012.
- L. Fillatre and I. Nikiforov. Non-bayesian detection and detectability of anomalies from a few noisy tomographic projections. *Signal Processing, IEEE Transactions on*, 55(2):401–413, 2007.
- M. Fouladirad and I. Nikiforov. Optimal statistical fault detection with nuisance parameters. *Automatica*, 41(7):1157–1171, 2005.
- M. Fouladirad, L. Freitag, and I. Nikiforov. Optimal fault detection with nuisance parameters and a general covariance matrix. *International Journal of Adaptive Control and Signal Processing*, 22(5):431–439, 2008.
- C.D. Godsil, G. Royle, and CD Godsil. *Algebraic graph theory*, volume 8. Springer New York, 2001.
- L. Hagen and A.B. Kahng. New spectral methods for ratio cut partitioning and clustering. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 11(9):1074–1085, 1992.
- Y.I. Ingster and I.A. Suslina. *Nonparametric goodness-of-fit testing under Gaussian models*, volume 169. Springer Verlag, 2003.
- L. Jacob, P. Neuvial, and S. Dudoit. Gains in power from structured two-sample tests of means on graphs. *Arxiv preprint arXiv:1009.5173*, 2010.
- M. Ledoux. *The concentration of measure phenomenon*, volume 89. Amer Mathematical Society, 2001.
- E.L. Lehmann and J.P. Romano. *Testing statistical hypotheses*. Springer Verlag, 2005.
- T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *Foundations of Computer Science, 1988., 29th Annual Symposium on*, pages 422–431. IEEE, 1988.

- J. Leskovec and C. Faloutsos. Scalable modeling of real graphs using kronecker multiplication. In *Proceedings of the 24th international conference on Machine learning*, pages 497–504. ACM, 2007.
- J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *The Journal of Machine Learning Research*, 11:985–1042, 2010.
- D.W. Matula and F. Shahrokhi. Sparsest cuts and bottlenecks in graphs. *Discrete Applied Mathematics*, 27(1):113–123, 1990.
- R. Merris. Laplacian graph eigenvectors. *Linear algebra and its applications*, 278(1):221–236, 1998.
- S. Milgram. The small world problem. *Psychology today*, 2(1):60–67, 1967.
- J.J. Moliterno, M. Neumann, and B.L. Shader. Tight bounds on the algebraic connectivity of a balanced binary tree. *Electronic Journal of Linear Algebra*, 6:62–71, 2000.
- O. Rojo. The spectrum of the laplacian matrix of a balanced binary tree. *Linear algebra and its applications*, 349(1):203–219, 2002.
- O. Rojo and R. Soto. The spectra of the adjacency matrix and laplacian matrix for some balanced trees. *Linear algebra and its applications*, 403:97–117, 2005.
- L. L. Scharf and B. Friedlander. Matched sub-space detectors. *Signal Processing, IEEE Transactions on*, 42(8):2146–2157, 1994.
- J. Sharpnack and A. Singh. Identifying graph-structured activation patterns in networks. In *Proceedings of Neural Information Processing Systems, NIPS*, 2010.
- J. Sharpnack, A. Rinaldo, and A. Singh. Sparsistency of the edge lasso over graphs. *AISats (JMLR WCP)*, 22:1028–1036, 2012.
- D.B. Shmoys. Cut problems and their application to divide-and-conquer. *Approximation algorithms for NP-hard problems*, pages 192–235, 1997.
- M. Talagrand. *The generic chaining*. Springer, 2005.
- A. Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of American Mathematical Society*, 54:426–482, 1943.

6 Appendix

6.1 Proofs in Section 2

Proof of Lemma 2. To expedite the proof, we express the LR statistics in terms of the sufficient statistics $\mathbf{y}_0 = \frac{1}{|\bar{C}|} \sum_{i \in C} \mathbf{y}_i \sim N(\beta_0, \sigma_0^2)$ and $\mathbf{y}_1 = \frac{1}{|C|} \sum_{i \in \bar{C}} \mathbf{y}_i \sim N(\beta_1, \sigma_1^2)$ for $\sigma_0 = \sigma/\sqrt{|\bar{C}|}$ and $\sigma_1 = \sigma/\sqrt{|C|}$. Then, we obtain

$$2 \log \Lambda_C(\mathbf{y}) = \frac{1}{\sigma_0^2} (\mathbf{y}_0 - \hat{\beta})^2 + \frac{1}{\sigma_1^2} (\mathbf{y}_1 - \hat{\beta})^2$$

where $\hat{\beta} = \frac{\sigma_1^2}{\sigma_0^2 + \sigma_1^2} \mathbf{y}_0 + \frac{\sigma_0^2}{\sigma_0^2 + \sigma_1^2} \mathbf{y}_1$ is the MLE under H_0 . (The likelihood under the alternative balances with the normalizing constant of the null likelihood.) Thus,

$$\begin{aligned} 2 \log \Lambda_C(\mathbf{y}) &= \frac{1}{\sigma_0^2} \left(\frac{\sigma_0^2}{\sigma_0^2 + \sigma_1^2} (\mathbf{y}_0 - \mathbf{y}_1) \right)^2 + \frac{1}{\sigma_1^2} \left(\frac{\sigma_1^2}{\sigma_0^2 + \sigma_1^2} (\mathbf{y}_0 - \mathbf{y}_1) \right)^2 \\ &= \frac{(\mathbf{y}_0 - \mathbf{y}_1)^2}{\sigma_0^2 + \sigma_1^2} = \frac{1}{\sigma^2} \frac{|C||\bar{C}|}{|V|} (\mathbf{y}_0 - \mathbf{y}_1)^2 \\ &= \frac{1}{\sigma^2} \frac{|V|}{|C||\bar{C}|} \left(\frac{|\bar{C}|}{|V|} \sum_{v \in C} \mathbf{y}_v - \frac{|C|}{|V|} \sum_{v \in \bar{C}} \mathbf{y}_v \right)^2 \\ &= \frac{1}{\sigma^2} \frac{|V|}{|C||\bar{C}|} \left(\sum_{v \in C} \mathbf{y}_v - \frac{|C|}{|V|} \sum_{v \in V} \mathbf{y}_v \right)^2 = \frac{1}{\sigma^2} \frac{|V|}{|C||\bar{C}|} \left(\sum_{v \in C} \tilde{\mathbf{y}}_v \right)^2. \end{aligned} \quad (11)$$

Now we let $\mathbf{x} = \mathbf{1}_C$, making the statistic above

$$2\sigma^2 \log \Lambda_C(\mathbf{y}) = \frac{\mathbf{x}^\top \tilde{\mathbf{y}} \tilde{\mathbf{y}} \mathbf{x}}{\mathbf{x}^\top \mathbf{K} \mathbf{x}} \text{ and } \frac{|\partial C||V|}{|C||\bar{C}|} = \frac{\mathbf{x}^\top \mathbf{L} \mathbf{x}}{\mathbf{x}^\top \mathbf{K} \mathbf{x}}.$$

The result now follows by considering all the indicator functions corresponding to the sets in \mathcal{C} . \square

6.2 Proofs in Section 3

Proof of Theorem 5. Let the true $C \in \mathcal{C}$ be known. The performance of the optimal test with C known, which by the Neyman-Pearson Lemma is based on $2 \log \Lambda_C(\mathbf{y})$, bounds the performance of that with C unknown. To this end, note that, under H_0 , the LR statistic (6) has a χ_1^2 , while under the alternative H_1^C it has a $\chi_1^2(\lambda)$ distribution with non-centrality parameter

$$\lambda = \frac{\delta^2}{\sigma^2} \frac{|C||\bar{C}|}{|V|} = \frac{\eta^2}{\sigma^2},$$

which is the square of the SNR. For fixed C , asymptotically indistinguishable of H_0 versus H_C^1 follows by considering any threshold and noticing that the associated type 1 and type 2 errors are non-vanishing under the SNR scaling assumed in the statement. Since the risk of testing H_0 versus H_1 is no smaller than the risk of testing H_0 versus H_C^1 , the result follows. \square

We remark that the proof of the previous result shows that when distinguishing H_0 from H_1^C , the power of the test is maximal when $|C| = |\bar{C}|$ for a fixed value of the SNR.

Proof of Lemma 7. Without loss of generality, let $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. We recall that, since G is connected, the combinatorial Laplacian \mathbf{L} is symmetric, its smallest eigenvalue is zero and the remaining eigenvalues are positive. By the spectral theorem, we can write $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$, where $\mathbf{\Lambda}$ is a $(n-1) \times (n-1)$ diagonal matrix containing the positive eigenvalues of \mathbf{L} in increasing order and the columns of the $n \times (n-1)$ matrix \mathbf{U} are the associated eigenvectors. Then, since each vector $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{1}^\top \mathbf{x} = 0$ can be written as $\mathbf{U} \mathbf{z}$ for a unique vector $\mathbf{z} \in \mathbb{R}^{n-1}$, we have

$$\begin{aligned} \mathcal{X} &= \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}^\top \mathbf{L} \mathbf{x} \leq \rho, \mathbf{x}^\top \mathbf{x} = 1, \mathbf{1}^\top \mathbf{x} \leq 0\} \\ &= \{\mathbf{U} \mathbf{z} \in \mathbb{R}^n : \mathbf{z} \in \mathbb{R}^{n-1}, \mathbf{z}^\top \mathbf{U}^\top \mathbf{L} \mathbf{U} \mathbf{z} \leq \rho, \mathbf{z}^\top \mathbf{U}^\top \mathbf{U} \mathbf{z} \leq 1\} \\ &= \{\mathbf{U} \mathbf{z} \in \mathbb{R}^n : \mathbf{z} \in \mathbb{R}^{n-1}, \frac{1}{\rho} \mathbf{z}^\top \mathbf{\Lambda} \mathbf{z} \leq 1, \mathbf{z}^\top \mathbf{z} \leq 1\}, \end{aligned}$$

where in the third identity we have used the fact that $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_{n-1}$. Letting $\mathcal{Z} = \{\mathbf{z} \in \mathbb{R}^{n-1} : \frac{1}{\rho} \mathbf{z}^\top \mathbf{\Lambda} \mathbf{z} \leq 1, \mathbf{z}^\top \mathbf{z} \leq 1\}$, we see that

$$\sup_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \mathbf{y} = \sup_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^\top \mathbf{U}^\top \mathbf{y} \stackrel{d}{=} \sup_{\mathbf{z} \in \mathcal{Z}} \mathbf{z}^\top \boldsymbol{\xi},$$

where $\boldsymbol{\xi} \sim N(0, \mathbf{I}_{n-1})$ and $\stackrel{d}{=}$ denotes equality in distribution.

Next, we show that the set \mathcal{Z} , which is the intersection of an ellipsoid with the unit ball in \mathbb{R}^{n-1} , is contained in an enlarged ellipsoid. The supremum of the Gaussian process $\mathbf{z}^\top \boldsymbol{\xi}$ over \mathcal{Z} will then be bounded by the supremum of the same process over this larger but simpler set, which we will be able to bound using directly a result from Talagrand [2005] based on chaining. To this end, let $\mathbf{A} = \frac{1}{\rho} \mathbf{\Lambda} = \text{diag}\{a_i\}_{i=1}^{n-1}$ and $d = \max\{j : a_j < 1\}$. For a vector $\mathbf{z} \in \mathbb{R}^{n-1}$ set $\mathbf{z}_1 = \mathbf{z}_{[d]}$, $\mathbf{z}_2 = \mathbf{z}_{[n-1] \setminus [d]}$, and $\mathbf{A}_2 = \text{diag}\{a_i\}_{i>d}$. Then, we observe the following chain of implications, holding for vectors $\mathbf{z} \in \mathbb{R}^{n-1}$:

$$\begin{aligned} \|\mathbf{z}\| \leq 1, \mathbf{z}^\top \mathbf{A} \mathbf{z} \leq 1 &\Rightarrow \|\mathbf{z}_1\| \leq 1, \sum_{i>d} a_i \mathbf{z}_i^2 \leq 1 \\ &\Rightarrow \mathbf{z}_1^\top \mathbf{z}_1 + \mathbf{z}_2^\top \mathbf{A}_2 \mathbf{z}_2 \leq 2 \Rightarrow \sum_i \frac{\max\{1, a_i\}}{2} \mathbf{z}_i^2 \leq 1. \end{aligned}$$

Hence, we have the bound

$$\mathbb{E} \sqrt{\hat{s}} \leq \mathbb{E} \sup_{\mathbf{z} \in \mathbb{R}^{n-1}} \mathbf{z}^\top \boldsymbol{\xi} \text{ s.t. } \sum_i 2 \max\{1, a_i\} \mathbf{z}_i^2 \leq 1.$$

Recalling that $a_i = \frac{\lambda_{i+1}}{\rho}$, for $i = 1, \dots, n-1$, where λ_{i+1} is the $(i+1)$ th eigenvalue of \mathbf{L} , by Proposition 2.2.1 in Talagrand [2005] the right hand side of the previous expression is bounded by $\sqrt{2 \sum_{i>1} \min\{1, \rho \lambda_i^{-1}\}}$. \square

Supplement to the proof of Theorem 6. The following property of Gaussian processes effectively reduces the study of their supremum to the study of its expectation. It was established by Borell [1975] and Cirelson et al. [1976] and can be found in Ledoux [2001].

Lemma 14. Consider a Gaussian process $\{Z_t\}_{t \in \mathcal{U}}$ where \mathcal{U} is compact with respect to metric

$$d(s, t) = (\mathbb{E}(Z_s - Z_t)^2)^{1/2}, \quad s, t, \in \mathcal{U},$$

and let $\sigma^2 \geq \sup_{t \in \mathcal{U}} \mathbb{E} Z_t^2$. We have that with probability at least $1 - \delta$

$$\left| \sup_{t \in \mathcal{U}} Z_t - \mathbb{E} \sup_{t \in \mathcal{U}} Z_t \right| < \sqrt{2\sigma^2 \log \frac{2}{\delta}}.$$

Notice that the natural distance is given by $d(\mathbf{x}_0, \mathbf{x}_1) = (\mathbb{E}((\mathbf{x}_0 - \mathbf{x}_1)^\top \mathbf{y})^2)^{1/2} = \sigma \|\mathbf{x}_0 - \mathbf{x}_1\|$ for $\mathbf{x}_0, \mathbf{x}_1 \in \mathcal{X}$. \square

Proof of Theorem 9. Recall that $\tilde{\mathbf{y}} = \mathbf{K} \mathbf{y}$, where $\mathbf{K} = \mathbf{I}_n - \frac{1}{n} \mathbf{1} \mathbf{1}^\top$ is the orthogonal projection matrix into the $(n-1)$ -dimensional linear subspace of vectors orthogonal to $\mathbf{1}$. Under H_0 , $\tilde{\mathbf{y}} \sim N(0, \sigma^2 \mathbf{K})$, and, therefore, $\|\tilde{\mathbf{y}}\|^2 \sim \chi_{n-1}^2$, since $\text{tr}(\mathbf{K}) = n-1$. On the other hand, under H_1^C for a fixed C , $\tilde{\mathbf{y}} \sim N(\mathbf{K} \boldsymbol{\beta}, \sigma^2 \mathbf{K})$, where $\boldsymbol{\beta}$ is given in as in (2). Thus, under H_1^C , $\|\tilde{\mathbf{y}}\|^2 \sim \chi_{n-1}^2(\lambda)$, where the non-centrality parameter is given by

$$\lambda = \boldsymbol{\beta}^\top \mathbf{K} \boldsymbol{\beta} = \frac{1}{2} \boldsymbol{\beta}^\top \mathbf{K}^\top \mathbf{K} \boldsymbol{\beta} = \|\boldsymbol{\beta} - \bar{\boldsymbol{\beta}}\|^2 \geq \frac{\eta^2}{2}, \quad (12)$$

where the second identity is due to the fact that \mathbf{K} is symmetric and idempotent and the last inequality to our assumption on the minimal separation η between H_0 and any of the alternatives. Thus, if $\eta/\sigma = \omega(\sqrt{n-1})$, then $\lambda = \omega(n-1)$. Hence, using standard chi-square tail bounds (see for example proposition 2 of Azizyan and Singh [2012]) and since the bound (12) holds uniformly over all $C \in \mathcal{C}$, it follows that the null and alternate are asymptotically distinguishable using the test statistic $\|\tilde{\mathbf{y}}\|$ if and only if $\frac{\eta}{\sigma} = \omega(\sqrt{n-1})$. \square

6.3 Proof in Section 4

Proof of Corollary 11. The study of the spectra of trees really began in earnest with the work of Fiedler [1975]. Notably, it became apparent that trees have eigenvalues with high multiplicities, particularly the eigenvalue 1. Molitierno et al. [2000] gave a tight bound on the algebraic connectivity of balanced binary trees (BBT). They found that for a BBT of depth ℓ , the reciprocal of the smallest eigenvalue ($\lambda_2^{(\ell)}$) is

$$\frac{1}{\lambda_2^{(\ell)}} \leq 2^\ell - 2\ell + 2 - \frac{2^\ell - \sqrt{2}(2\ell - 1 - 2^{\ell-1})}{2^\ell - 1 - \sqrt{2}(2^{\ell-1} - 1)} + (3 - 2\sqrt{2} \cos(\frac{\pi}{2\ell - 1}))^{-1} \leq 2^\ell + 105I\{\ell < 4\} \quad (13)$$

Rojo [2002] gave a more exact characterization of the spectrum of a balanced binary tree, providing a decomposition of the Laplacian's characteristic polynomial. Specifically, the characteristic polynomial of \mathbf{L} is given by

$$\det(\lambda \mathbf{I} - \mathbf{L}) = p_1^{2^{\ell-2}}(\lambda) p_2^{2^{\ell-3}}(\lambda) \dots p_{\ell-3}^{2^2}(\lambda) p_{\ell-2}^2(\lambda) p_{\ell-1}(\lambda) s_\ell(\lambda) \quad (14)$$

where $s_\ell(\lambda)$ is a polynomial of degree ℓ and $p_i(\lambda)$ are polynomials of degree i with the smallest root satisfying the bound in (13) with ℓ replaced with i . In Rojo and Soto [2005], they extended this work to more general balanced trees.

By (14) we know that at most $\ell + (\ell - 1) + (\ell - 2)2 + \dots + (\ell - j)2^{j-1} \leq \ell 2^j$ eigenvalues have reciprocals larger than $2^{\ell-j} + 105I\{j < 4\}$. Let $k = \max\{\lceil \frac{\ell}{e} 2^{\ell(1-\alpha)} \rceil, 2^3\}$, then we have ensured that at most k eigenvalues are smaller than ρ . For n large enough

$$\sum_{i>1} \min\{1, \rho \lambda_i^{-1}\} \leq k + \rho \sum_{j>\log k}^{\ell} \ell 2^j 2^{\ell-j} = k + \ell(\ell - \log k) n \rho = O(n^{1-\alpha} (\log n)^2)$$

□

Proof of Corollary 13. The Kronecker product of two matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ is defined as $\mathbf{A} \otimes \mathbf{B} \in \mathbb{R}^{(n \times n) \times (n \times n)}$ such that $(\mathbf{A} \otimes \mathbf{B})_{(i_1, i_2), (j_1, j_2)} = A_{i_1, j_1} B_{i_2, j_2}$. Some matrix algebra shows that if H_1 and H_2 are graphs on p vertices with Laplacians $\mathbf{L}_1, \mathbf{L}_2$ then the Laplacian of their Kronecker product, $H_1 \otimes H_2$, is given by $\mathbf{L} = \mathbf{L}_1 \otimes \mathbf{I}_p + \mathbf{I}_p \otimes \mathbf{L}_2$ (Merris [1998]). Hence, if $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^p$ are eigenvectors, viz. $\mathbf{L}_1 \mathbf{v}_1 = \lambda_1 \mathbf{v}_1$ and $\mathbf{L}_2 \mathbf{v}_2 = \lambda_2 \mathbf{v}_2$, then $\mathbf{L}(\mathbf{v}_1 \otimes \mathbf{v}_2) = (\lambda_1 + \lambda_2) \mathbf{v}_1 \otimes \mathbf{v}_2$, where $\mathbf{v}_1 \otimes \mathbf{v}_2$ is the usual tensor product. This completely characterizes the spectrum of Kronecker products of graphs.

We should argue the choice of $\rho \propto p^{2k-\ell-1}$, by showing that it is the results of cuts at level k . We say that an edge $e = ((i_1, \dots, i_\ell), (j_1, \dots, j_\ell))$ has scale k if $i_k \neq j_k$. Furthermore, a cut has scale k if each of its constituent edges has scale at least k . Each edge at scale k has weight $p^{k-\ell}$ and there are $p^{\ell-1}$ such edges, so cuts at scale k have total edge weight bounded by

$$p^{\ell-1} \sum_{i=1}^k p^{i-\ell} = p^{k-1} \frac{p - \frac{1}{p^{k-1}}}{p - 1} \leq \frac{p^k}{p - 1}$$

Cuts at scale k leave components of size $p^{\ell-k}$ intact, meaning that $\rho \propto p^{2k-\ell-1}$ for large enough p .

We now control the spectrum of the Kronecker graph. Let the eigenvalues of the base graph H be $\{\nu_j\}_{j=1}^p$ in increasing order. The eigenvalues of G are precisely the sums

$$\lambda_i = \frac{1}{p^{\ell-1}} \nu_{i_1} + \frac{1}{p^{\ell-2}} \nu_{i_2} + \dots + \frac{1}{p} \nu_{i_{\ell-1}} + \nu_{i_\ell}$$

for $i = (i_j)_{j=1}^\ell \subseteq [p]$. The eigenvalue distribution $\{\lambda_i\}$ stochastically bounds

$$\lambda_i \geq \sum_{j=1}^{\ell} \frac{1}{p^{\ell-j}} \nu_{i_j} I\{\nu_{i_j} \neq 0\} \geq \frac{\nu_2}{p^{Z(i)}}$$

where $Z(i) = \min\{j : \nu_{i_{\ell-j}} \neq 0\}$. Notice that if i is chosen uniformly at random then $Z(i)$ has a geometric distribution with probability of success $(p - 1)/p$. Also $\rho / (\frac{\nu_2}{p^{Z(i)}}) = p^{Z(i)+2k-\ell-1} / \nu_2 \geq 1$ if $Z(i) \geq \ell + 1 - 2k +$

$\log_p \nu_2$, so

$$\frac{1}{p^\ell} \sum_{i \in [p]^\ell} \min\{1, \frac{\rho}{\lambda_i}\} \leq \frac{p^{2k-\ell-1}}{\nu_2} + \sum_{Z=1}^{\lfloor \ell+1-2k+\log_p \nu_2 \rfloor} \frac{p^{Z+2k-\ell-1}}{\nu_2} \frac{1}{p^Z} \frac{p-1}{p} \leq \frac{(\ell+2)p^{2k-\ell-1}}{\nu_2}$$

This followed from the geometric probability mass function. We also know that the algebraic connectivity, ν_2 , is bounded from below by $4p^{-2}$, so the following result holds. □